Bayesian Nonparametric Dimensionality Reduction of Categorical Data for Predicting Severity of COVID-19 in Pregnant Women

Marzieh Ajirak,[°] Cassandra Heiselman,^{*} Anna Fuchs,^{*} Mia Heiligenstein,^{*} Kimberly Herrera,^{*} Diana Garretto,^{*} and Petar M. Djurić,[°] [°]Department of Electrical and Computer Engineering, Stony Brook University ^{*}Department of Obstetrics, Gynecology and Reproductive Medicine

Stony Brook, NY 11794, USA

Abstract—The coronavirus disease (COVID-19) has rapidly spread throughout the world and while pregnant women present the same adverse outcome rates, they are underrepresented in clinical research. We collected clinical data of 155 test-positive COVID-19 pregnant women at Stony Brook University Hospital. Many of these collected data are of multivariate categorical type, where the number of possible outcomes grows exponentially as the dimension of data increases. We modeled the data within the unsupervised Bayesian framework and mapped them into a lower dimensional space using latent Gaussian processes. The latent features in the lower dimensional space were further used for predicting if a pregnant woman would be admitted to a hospital due to COVID-19 or would remain with mild symptoms. We compared the prediction accuracy with the dummy/one-hot encoding of categorical data and found that the latent Gaussian process had better accuracy.

Index Terms—Categorical Latent Gaussian Process, Coronavirus Disease, Data mining

I. INTRODUCTION

The coronavirus disease 2019 (COVID-19) has become an unprecedented public health crisis. Around the world, many governments issued a call to researchers in machine learning (ML) and artificial intelligence (AI) to address highpriority questions related to COVID-19. This call was not unusual because ML methods are finding many uses in medical diagnosis applications. The ML field is rich with examples where based on predictive models one can estimate disease severity [1] and consequently, the state of a patient's health [2]–[4]. These models employ data-driven algorithms that can extract features and discover complicated patterns that could have not been recognized or interpreted by humans.

Pregnant women are a particularly important patient population to study due to their vulnerability to disease and the often underrepresentation of the population in clinical research [5]. Despite studies in this field [6]–[10], there has been a relative sparsity of data in regards to COVID-19 and its effect on pregnancy. Utilizing ML techniques to study this population during the pandemic can help build pregnancyspecific evidence to guide clinical recommendations [11]. Much of the medical data are of multivariate categorical type, and typically they represent patients' demographics, maternal comorbidities, pregnancy complications, and disease symptoms. As a result, one has to work with long vectors of categorical variables which in turn leads to a huge number of possible realizations. This then creates very sparse spaces when we deal with a limited number of data [12]. Besides, the data include random errors and systematic biases, and sometimes they are missing [13]. By overcoming the challenges that clinical data introduce, one can layout the grounds for developing more accurate models and efficient algorithms for inference.

The key to have successful predictive methods largely depends on feature selection and data representation. A common approach is to have a clinical doctor specify the variables and label the clinical data to be used as training sets. Then the ML method will find mappings and features from the data, which subsequently will be tested on new data sets. Although appropriate in many situations, a supervised definition of the features contributes to losing an opportunity to learn latent patterns and features [4]. In countering the subjectivity of defining the features, an unsupervised learning approach can be used to extract useful information from data. One other advantage of unsupervised learning is that abstract features of patients can often be represented in low-dimensional spaces and thus, they can summarize efficiently the information available in the data. This further allows for easy visualization of the cohort of patients under consideration.

In the ML literature, categorical latent Gaussian processes provide data efficient and powerful Bayesian framework for learning latent functions or patterns [14]. In this paper, we model the categorical data from pregnant women as generated non-linearly from a latent space. More specifically, we map the categorical variables including maternal comorbidities, pregnancy complications, ABO blood types, etc., into a continuous lower dimensional space. Then we use these learned features along with the remaining numerical data (maternal age, BMI, etc.) to predict whether (a) the patient will develop severe symptoms and will come back to the hospital due to COVID-19, days after tested positive, or (b) the patient will remain

The authors thank the support of NIH under Award RO1HD097188-01.

 Table I

 LIST OF SYMBOLS AND NOTATIONS

N	_	Number of patients
D	_	Number of categorical variables
K	-	Number of possible outcomes
Q	-	Dimensionality of the latent space
$oldsymbol{x}_n$	-	Latent value of the <i>n</i> th patient
$oldsymbol{y}_n$	-	A D-dimensional observed vector
y_{nd}	-	The <i>d</i> th element of \boldsymbol{y}_n that can take one of <i>D</i>
		categorical values
$oldsymbol{f}_{nd}$	-	Vector of probability weights of the dth cate-
		gorical variable for the <i>n</i> th patient
\mathbf{K}_d	-	Covariance matrix corresponding to the dth
		categorical variables
κ	-	A kernel function
${\mathcal F}$	-	A sample function from $\mathcal{GP}(.,.)$
Z	-	Inducing input locations
U	-	Inducing variables

asymptomatic or symptomatic but with mild symptoms. We compared the performance obtained by direct and non-linear dimensionality reduction of the categorical data with the methodology of one-hot encoding, which is commonly applied in the machine learning circles when dealing with categorical data.

The remainder of this paper is organized as follows. In the next section, we explain the categorical latent Gaussian model first introduced in [14]. Then we introduce an alternative pipeline that deals with categorical data. We test the proposed approach first on synthesized data and then on original COVID-19 data.

II. BACKGROUND ON GPLVM

Gaussian process latent variable models (GPLVMs) are Bayesian nonparametric frameworks that allow for unsupervised learning [15]. GPLVMs can be seen as multi-output Gaussian process regressions when the inputs are unobserved. To be more specific, let $f_n \in \mathbb{R}^K$ be the *n*-th observed data vector of dimension K, n = 1, 2, ..., N. Further, let these data be associated with inputs $x_n \in \mathbb{R}^Q$ through K different functions. If we assume that these functions are independent, then for f_n we can write

$$p(\boldsymbol{f}_n(\boldsymbol{x}_n)) = \prod_{k=1}^{K} p(f_{nk}(\boldsymbol{x}_n)), \qquad (1)$$

where $f_{nk}(\boldsymbol{x}_n)$ represents the kth dimension of $f_n(\boldsymbol{x}_n)$ and

$$p(f_{nk}(\boldsymbol{x}_n)) = \mathcal{N}(f_{nk}; 0, k(\boldsymbol{x}_n, \boldsymbol{x}'_n)), \qquad (2)$$

where the notation $\mathcal{N}(f_{nk}; 0, k(\boldsymbol{x}_n, \boldsymbol{x}'_n))$ means that the random variable $f_{nk}(\boldsymbol{x}_n)$ is Gaussian with mean zero and variance defined by the covariance function $k(\boldsymbol{x}_n, \boldsymbol{x}'_n)$. In order to automatically learn the dimensionality of the latent space, we will use the concept known as Automatic Relevance Determination (ARD) with the kernel

$$\kappa\left(\boldsymbol{x}_{n},\boldsymbol{x}_{n}'\right) = \sigma_{f}^{2} \exp\left(-\frac{1}{2} \sum_{q=1}^{Q} \alpha_{q} \left(x_{n,q} - x_{n,q}'\right)^{2}\right). \quad (3)$$

In GPLVMs, $X \in \mathbb{R}^{N \times Q}$ is a matrix of latent variables, and therefore we assign it a prior density. A typical approach is to use the standard Gaussian distribution, and thus we have

$$p(\boldsymbol{X}) = \prod_{n=1}^{N} \mathcal{N}(\boldsymbol{x}_n; \boldsymbol{0}, \boldsymbol{I}_Q), \qquad (4)$$

where the x_n 's are the rows of X. By defining the matrix of observations $F \in \mathbb{R}^{N \times K}$, where the rows represent the multiple outputs f_n , we wish to compute the marginal likelihood of the data:

$$p(\mathbf{F}) = \int p(\mathbf{F}|\mathbf{X}) p(\mathbf{X}) d\mathbf{X}.$$
 (5)

The authors in [15] developed a variational Bayesian approach for the marginalization of the latent variables, X, allowing them to optimize the resulting lower bound on the marginal likelihood with respect to the hyperparameters. They further used the lower bound for model comparison and automatic selection of the latent dimensionality.

III. MULTIVARIATE DISCRETE GPLVM

A. Generative Model

We consider now the discrete version of GPLVM where for each input x_n , we observe a discrete variable y_n that can take values 1, ..., K, with probabilities

$$p(y_n = k) = \frac{\exp(f_{nk})}{\sum_{k'=1}^{K} \exp(f_{nk'})}.$$
 (6)

In the multivariate case, we have $y_n \in \mathbb{R}^D$. Next, we consider a generative model for a dataset $Y \in \mathbb{R}^{N \times D}$ with N observations and D categorical variables. We denote the d-th variable in the n-th observation by y_{nd} . Now we express (6) as

$$p(y_{nd} = k) = \frac{\exp(f_{ndk})}{\sum_{k'=1}^{K} \exp(f_{ndk'})},$$
(7)

where f_{ndk} is function of the input variable $x_n \in \mathbb{R}^Q$, i.e., $f_{ndk} = \mathcal{F}_{dk}(x_n)$. Next, we summarize the generative model (the indices below have the following meaning: n refers to observation, d to the dimension of the output, m to an inducing point (defined below), and k to category),

$$x_{nq} \stackrel{\text{id}}{\sim} \mathcal{N}\left(0, \sigma_x^2\right),$$
 (8)

$$\mathcal{F}_{dk} \stackrel{\text{nd}}{\sim} \mathcal{GP}\left(0, k_d(\cdot, \cdot)\right), \tag{9}$$

$$f_{ndk} = \mathcal{F}_{dk}\left(\boldsymbol{x}_n\right),\tag{10}$$

$$u_{mdk} = \mathcal{F}_{dk}\left(\boldsymbol{z}_{m}\right),\tag{11}$$

$$p(y_{nd} = k) = \frac{\exp(f_{ndk})}{\sum_{k'=1}^{K} \exp(f_{ndk'})},$$
 (12)

where x_{nq} and \mathcal{F}_{dk} are latent variables with prior distributions given by (8) and (9), respectively, with \mathcal{GP} signifying Gaussian process (GP). Further, the z_m s are inducing inputs, $m = 1, 2, \ldots, M$, and the u_{mdk} s are inducing outputs whose role is explained further below. We note that we assume a Gaussian distribution prior with standard deviation σ_x^2 for x_{nk} , and a GP prior for each of the functions \mathcal{F} . We reiterate that for each vector of latent function values f_{dk} , we introduce a separate set of M variational inducing variables u_{dk} , evaluated at a set of inducing input locations from the set $\mathcal{Z} = \{z_1, z_2, \ldots, z_M\}$. It is assumed that all u_{dk} s are computed at the same inducing locations. The inducing variables are function points drawn from the GP prior and lie in the same latent space as F variables (Fig. 1). The pictorial description of the generative model is displayed in Fig. 2.



Figure 1. Latent weights and inducing variables.



Figure 2. A graphical representation of the generative model.

B. Inference

The marginal log-likelihood is intractable because of the covariance function of the GP and the nonlinear Softmax likelihood in (12). We consider a variational approximation to the posterior distribution of $\boldsymbol{X}, \boldsymbol{F}$ and $\boldsymbol{U} \in \mathbb{R}^{M \times D \times K}$ factorized as,

$$q(\boldsymbol{X}, \boldsymbol{F}, \boldsymbol{U}) = q(\boldsymbol{X})q(\boldsymbol{U})p(\boldsymbol{F}|\boldsymbol{X}, \boldsymbol{U}). \quad (13)$$

By applying Jensen's inequality, we can write a lower bound of the log-evidence (ELBO) as

$$\log p(\boldsymbol{Y}) = \log \int p(\boldsymbol{X}) p(\boldsymbol{U}) p(\boldsymbol{F}|\boldsymbol{X}, \boldsymbol{U}) p(\boldsymbol{Y}|\boldsymbol{F}) d\boldsymbol{X} d\boldsymbol{F} d\boldsymbol{U}$$

$$\geq -\mathrm{KL}(q(\boldsymbol{X}) \| p(\boldsymbol{X})) - \mathrm{KL}(q(\boldsymbol{U}) \| p(\boldsymbol{U}))$$

$$+ \sum_{n=1}^{N} \sum_{d=1}^{D} \int q(\boldsymbol{x}_{n}) q(\boldsymbol{U}_{d}) p(\boldsymbol{f}_{nd} | \boldsymbol{x}_{n}, \boldsymbol{U}_{d})$$

$$\cdot \log p(\boldsymbol{y}_{nd} | \boldsymbol{f}_{nd}) d\boldsymbol{x}_{n} d\boldsymbol{f}_{nd} \boldsymbol{U}_{d} := \mathcal{L}, \quad (14)$$

where,

$$p(\boldsymbol{f}_{nd}|\boldsymbol{x}_n, \boldsymbol{U}_d) = \prod_{k=1}^{K} \mathcal{N}(f_{ndk}; \mathbf{k}_{d,nM}^{\top} \mathbf{K}_{d,MM}^{-1} \boldsymbol{u}_{dk}, \\ k_{d,nn} - \mathbf{k}_{d,nM}^{\top} \mathbf{K}_{d,MM}^{-1} \mathbf{k}_{d,Mn}).$$
(15)

The lower bound is still intractable because of the softmax likelihood, $\log p(\mathbf{y}_{nd} | \mathbf{f}_{nd})$. Therefore, we will compute the lower bound \mathcal{L} and its derivatives with the Monte Carlo method. We draw samples of $\mathbf{x}_n, \mathbf{U}_d \in \mathbb{R}^{M \times K}$ (see Fig. 1) and \mathbf{f}_{nd} from $q(\mathbf{x}_n), q(\mathbf{U}_d)$, and $p(\mathbf{f}_{nd} | \mathbf{x}_n, \mathbf{U}_d)$, respectively, and estimate \mathcal{L} with the sample average. We consider mean field variational approximation of the latent points $q(\mathbf{X})$ and a joint Gaussian distribution for $q(\mathbf{U})$ as,

$$q(\boldsymbol{U}) = \prod_{d=1}^{D} \prod_{k=1}^{K} \mathcal{N}\left(\boldsymbol{u}_{dk}; \boldsymbol{\mu}_{dk}, \boldsymbol{\Sigma}_{d}\right), \qquad (16)$$

$$q(\boldsymbol{X}) = \prod_{n=1}^{N} \prod_{q=1}^{Q} \mathcal{N}\left(x_{nq}; m_{nq}, \sigma_{nq}^{2}\right), \qquad (17)$$

where the covariance matrix Σ_d is shared for the same categorical variable *d*. The KL divergence in \mathcal{L} can be computed analytically with the given variational distributions. We need to optimize the hyperparameters of each GP (parameters of \mathbf{K}_d), parameters of the variational random variables u_{dk} , μ_{dk} , Σ_d , mean m_{nq} and variance σ_{nq}^2 of the latent inputs.

IV. EXPERIMENTS AND RESULTS

A. 1-D Input and Output

Consider the categorical variable y that can take values 001, 010, and 100 (or blue, red, and green). The input variable xcomes from a space of patients. Further, let three functions $\mathcal{F}_{11}, \mathcal{F}_{12}$, and \mathcal{F}_{13} model the f_{ndk} s that are used for computing the probability of each category (the first index of the functions refers to the dimension, which in this example equals to one). For instance, $\mathcal{F}_{11}(x_n)$ is proportional to the probability of yfor patient n with input x_n being 001. Similarly, we define f_{n12} and f_{n13} for the categories 010 and 100, respectively.

We perform the inference using the introduced method (Fig. 3 (a)) and compare it with the one-hot encoding of the categorical variables (Fig. 3 (b)). We observe that by the one-hot encoding and then applying GPLVM, the structure of the latent space is distorted. The first two dimensions of x are shown in Fig. 4 (c). Although a one-dimensional manifold is detected, the points at the boundary of the two clusters are obviously distorted.

B. COVID-19 Data

We used data collected at SBUH of 155 test-positive COVID-19 pregnant women. The dataset is composed of categorical variables including patients' symptoms, maternal comorbidities, pregnancy complications, race, employer type, insurance, known sick contact, and ABO blood type. It also has numerical data including age, BMI, gravidity, parity, and admission lab values. The list of categorical and numerical



(a) Embedding the categorical variables into lower dimension space.



(b) One-hot encoding of the categorical variables before dimensionality reduction.

Figure 3. Model Pipeline

Table II LIST OF COVID-19 PATIENTS VARIABLES

Input Variables	Туре		
Maternal age (year)	Numerical		
BMI (kg/m2)	Numerical		
Gravidity	Numerical		
Parity	Numerical		
Admission Lab Values	Numerical		
Thermodynamic Symptoms	Categorical		
Lower Respiratory Symptoms	Categorical		
HEENT Symptoms	Categorical		
GI symptoms	Categorical		
Hemodynamic Symptoms	Categorical		
Cardiovascular Symptoms	Categorical		
Musculoskeletal Symptoms	Categorical		
Race	Categorical		
Employer type	Categorical		
Insurance	Categorical		
Known sick contact type	Categorical		
Maternal comorbidities	Categorical		
Pregnancy complications	Categorical		
ABO blood type	Categorical		
Symptoms at time of diagnosis	Binary (Yes/No)		
Admitted to hospital for COVID-19	Binary (Yes/No)		
Admitted to ICU	Binary (Yes/No)		

variables is summarized in Table II. The cohort consisted of 60 asymptomatic cases, 81 moderate symptomatic, and 14 severely ill patients who were admitted to hospital for COVID-19. Of the latter 14 patients, four were admitted to ICU.

We first reduced the dimension of categorical data by mapping them into a lower-dimension space using discrete-GPLVM. Next, we used the extracted latent features combined with numerical variables for the supervised task of binary classification. Then we converted the categorical variables to one-hot features and then applied GPLVM. For classification we employed Random Forest, Naïve Bayes, AdaBoost, *k*-Nearest Neighbours (kNN), Support Vector Machine (SVM), and Logistic Regression. We compared the performances of the methods by Area Under the ROC Curve (AUC), Classifi-







cation Accuracy (CA), F1, Precision, and Recall, where

$$CA = \frac{TP + TN}{TP + TN + FP + FN},$$
(18)

with TP representing True Positive, TN True Negative, FP False Positive, and FN False Negative predictions,

$$\operatorname{Recall} = \frac{TP}{TP + FN},\tag{19}$$

$$Precision = \frac{TP}{TP + FP},$$
(20)

and

$$F1 = 2 * \frac{\text{Recall * Precision}}{\text{Recall + Precision}}.$$
 (21)

The results are summarized in Tables III and IV.

The results suggests that the performance of almost all classifiers improved by using the discrete GPLVM. The best

Table III DISCRETE-GPLVM

Model	AUC	CA	F1	Precision	Recall
Random Forest	0.842	0.787	0.788	0.789	0.787
Naïve Bayes	0.781	0.639	0.631	0.751	0.639
AdaBoost	0.715	0.723	0.724	0.728	0.723
kNN	0.674	0.671	0.655	0.661	0.671
SVM	0.670	0.600	0.604	0.613	0.600
Logistic Regression	0.601	0.619	0.611	0.608	0.619

Table IV ONE-HOT ENCODING-GPLVM

Model	AUC	CA	F1	Precision	Recall
Random Forest	0.770	0.729	0.729	0.729	0.729
SVM	0.702	0.677	0.642	0.678	0.677
Naïve Bayes	0.689	0.568	0.550	0.692	0.568
kNN	0.677	0.665	0.647	0.653	0.665
AdaBoost	0.671	0.684	0.685	0.687	0.684
Logistic Regression	0.607	0.606	0.601	0.598	0.606

performance of all classifiers was achieved by Random Forest. It appears that with dimensionality reduction using discrete GPLVM we compress information better than with GPLVM carried out by one-hot encoding.

We also mapped the data for the task of visualization of the cohort. Figure 5 shows the visualization of the patients using discrete-GPLVM by setting the latent dimension to Q = 2. We observe that the latent features of the symptomatic patients or patients with mild symptoms (blue circles) are well clustered and somewhat separated from the patients who were hospitalized or who were admitted to ICU (red circles and red crosses).



Figure 5. Visualization of the patients. Blue circles represent asymptomatic patients or patients with mild symptoms, red circles represent patients who were hospitalized and red crosses are patients who were admitted to ICU.

V. CONCLUSION

In this paper, we modeled multivariate categorical data using Gaussian process latent variable models to predict if a pregnant women would be admitted to the hospital due to COVID-19. In our approach, we used a data-efficient Bayesian framework for reducing the dimension of high-dimensional categorical data. Our tests with synthetic data showed that the method is capable of finding latent structures of the data. Further, the results on test-positive COVID-19 pregnant women suggest that the method discovered latent structures that were useful for further classification of the data.

REFERENCES

- X. Jiang, M. Coffee, A. Bari, J. Wang, X. Jiang, J. Huang, J. Shi, J. Dai, J. Cai, T. Zhang *et al.*, "Towards an artificial intelligence framework for data-driven prediction of coronavirus clinical severity," *Computers, Materials & Continua*, vol. 63, no. 1, pp. 537–551, 2020.
- [2] H. Yao, N. Zhang, R. Zhang, M. Duan, T. Xie, J. Pan, E. Peng, J. Huang, Y. Zhang, X. Xu *et al.*, "Severity detection for the coronavirus disease 2019 (covid-19) patients using a machine learning model based on the blood and urine tests," *Frontiers in Cell and Developmental Biology*, vol. 8, p. 683, 2020.
- [3] A. Alimadadi, S. Aryal, I. Manandhar, P. B. Munroe, B. Joe, and X. Cheng, "Artificial intelligence and machine learning to fight COVID-19," *Physiol Genomics*, vol. 52, pp. 200–202, 2020.
- [4] R. Miotto, L. Li, B. A. Kidd, and J. T. Dudley, "Deep patient: An unsupervised representation to predict the future of patients from the electronic health records," *Scientific Reports*, vol. 6, no. 1, pp. 1–10, 2016.
- [5] G. Favre, L. Pomar, and D. Baud, "Coronavirus disease 2019 during pregnancy: Do not underestimate the risk of maternal adverse outcomes," *American Journal of Obstetrics & Gynecology MFM*, vol. 2, no. 3, pp. 100160–100160, Aug 2020. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/32838265
- [6] N. Breslin, C. Baptiste, C. Gyamfi-Bannerman, R. Miller, R. Martinez, K. Bernstein, L. Ring, R. Landau, S. Purisch, A. M. Friedman *et al.*, "COVID-19 infection among asymptomatic and symptomatic pregnant women: Two weeks of confirmed presentations to an affiliated pair of new york city hospitals," *American Journal of Obstetrics & Gynecology MFM*, p. 100118, 2020.
- [7] Y. Wu, C. Zhang, H. Liu, C. Duan, C. Li, J. Fan, H. Li, L. Chen, H. Xu, X. Li *et al.*, "Perinatal depressive and anxiety symptoms of pregnant women during the coronavirus disease 2019 outbreak in china," *American Journal of Obstetrics and Gynecology*, vol. 223, no. 2, pp. 240–e1, 2020.
- [8] H. Chen, J. Guo, C. Wang, F. Luo, X. Yu, W. Zhang, J. Li, D. Zhao, D. Xu, Q. Gong *et al.*, "Clinical characteristics and intrauterine vertical transmission potential of COVID-19 infection in nine pregnant women: a retrospective review of medical records," *The Lancet*, vol. 395, no. 10226, pp. 809–815, 2020.
- [9] P. Dashraath, J. L. J. Wong, M. X. K. Lim, L. M. Lim, S. Li, A. Biswas, M. Choolani, C. Mattar, and L. L. Su, "Coronavirus disease 2019 (COVID-19) pandemic and pregnancy," *American Journal of Obstetrics* and Gynecology, vol. 222, no. 6, pp. 521–531, 2020.
- [10] R. A. Pierce-Williams, J. Burd, L. Felder, R. Khoury, P. S. Bernstein, K. Avila, C. A. Penfield, A. S. Roman, C. A. DeBolt, J. L. Stone *et al.*, "Clinical course of severe and critical coronavirus disease 2019 in hospitalized pregnancies: a united states cohort study," *American Journal of Obstetrics & Gynecology MFM*, vol. 2, no. 3, p. 100134, 2020.
- [11] X. Chen and Z. Liu, "Early prediction of mortality risk among severe COVID-19 patients using machine learning," *MedRxiv*, 2020.
- [12] A. Agresti, An Introduction to Categorical Data Analysis. John Wiley & Sons, 2018.
- [13] R. J. Little and D. B. Rubin, *Statistical analysis with missing data*. John Wiley & Sons, 2019, vol. 793.
- [14] Y. Gal, Y. Chen, and Z. Ghahramani, "Latent Gaussian processes for distribution estimation of multivariate categorical data," in *International Conference on Machine Learning*, 2015, pp. 645–654.
- [15] M. Titsias and N. D. Lawrence, "Bayesian Gaussian process latent variable model," in *Proceedings of the Thirteenth International Conference* on Artificial Intelligence and Statistics, 2010, pp. 844–851.